

Sebastian Krey

## HPC at GWDG

Site update and data protection on HPC systems

# GWDG

## Structure

- Joint data and IT-comptence center for the University of Göttingen, University clinic and the Max-Planck-Society
- Established in 1970
- Shareholders Max-Planck-Society and Unversity Göttingen, each 50%

## Mission

- It is the mission of GWDG to support and enable science – now and in future – by applying appropriate technologies, developing innovative solutions, and providing reliable services
- Plan, implement and operate IT infrastructures together with/for customers

## Central responsibilities

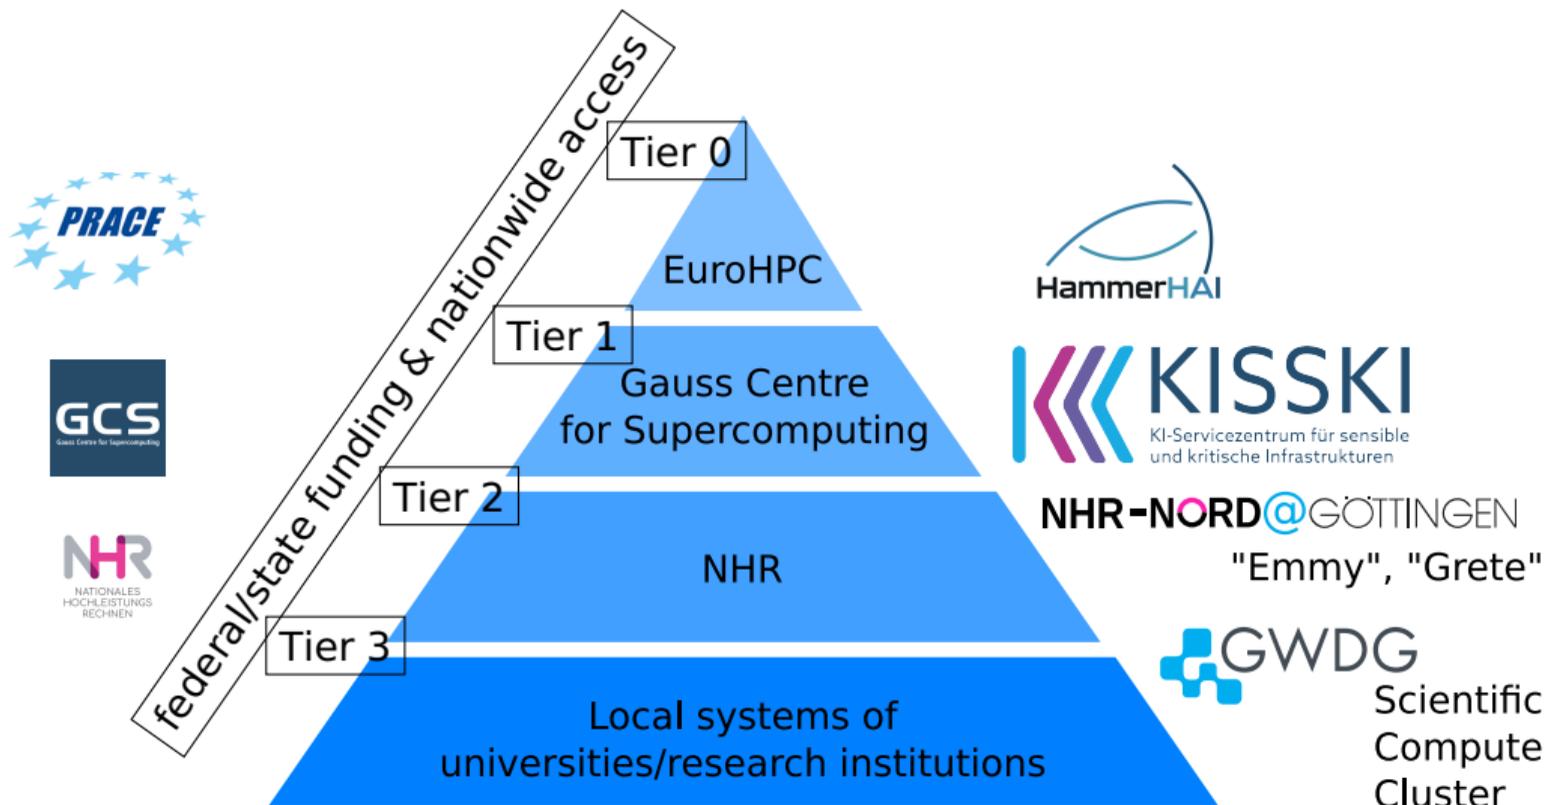
- Modern and secure IT infrastructure
- IT support for excellent research
- In-house research for innovative IT services

# GWDG

## Supra-regional tasks

- National High Performance Computing Center (NHR-NORD@Göttingen)
- National HPC Center of the DLR
- Upcoming regional HPC center for Lower Saxony (HPC.NDS)
- AI service center for sensitive and critical infrastructures (KISSKI)
- Partner in the AI Factory HammerHA1
- Data center in four NFDI consortia
- Host for DARIAH-EU, German National Library, GFBio, NUM CODEX, MWS, WirLernenOnline, Text+, etc.
- Cloud operator, including the Academic Cloud for universities in Germany

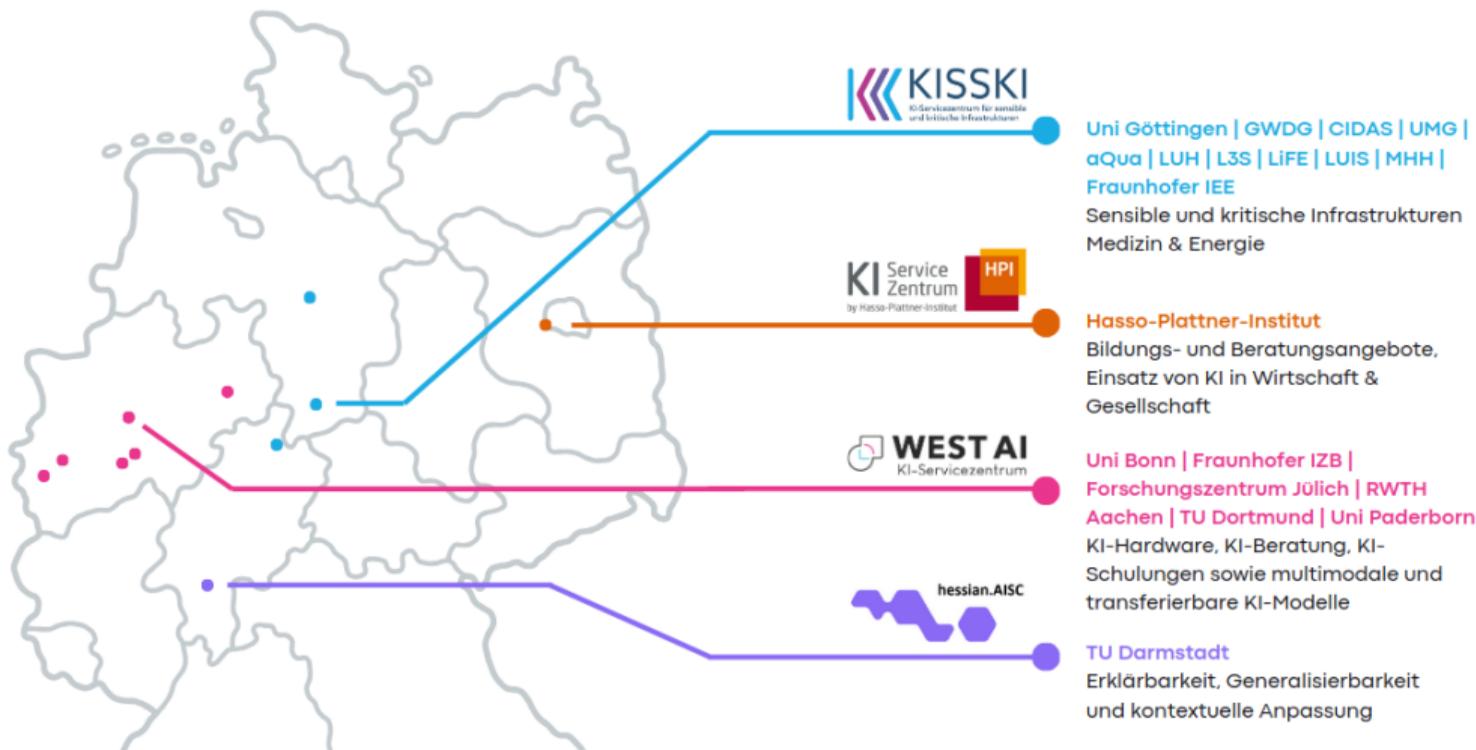
# German HPC Infrastructure



## NHR Alliance - Members

- Rhein-Westfälische Technische Hochschule Aachen
- Zuse Institute Berlin (ZIB)
- Technische Universität Darmstadt
- Technische Universität Dresden
- Friedrich-Alexander-Universität (FAU) Erlangen-Nürnberg
- GWDG/Georg-August-Universität Göttingen
- Karlsruhe Institute of Technology
- Paderborn University
- Consortium Süd-West (Goethe University Frankfurt, Rheinland-Pfälzische Technische Universität Kaiserslautern-Landau, Johannes Gutenberg Universität Mainz, Universität des Saarlandes)

# KISSKI - ein KI-Servicezentrum für Kritische Infrastrukturen



# Services for HPC and AI

- Infrastructure
- Scientific Software
- Consulting
- Training
- Support

One current focus is on usability, availability and security

# Our Offer - AI for Everyone

- Comprehensive offering around AI
- Solutions available in browser and via API
  - ▶ ChatAI - Chatbot
- Hardware resources
  - ▶ From ARM, RISC-V to Graphcore and more
- Consulting services
- Custom development projects
- Training programs in HPC and AI
  - ▶ GWDG Academy
  - ▶ KISSKI Training Offer
  - ▶ Custom courses available upon request



# CPU System “Emmy”

## ■ Phase 1 (2018) → out of operation

- ▶ 448 standard nodes (2 Xeon Gold 6148, 2x20 CPU-Cores, 192 GB memory)
- ▶ High Mem nodes: 16x 768 GB

## ■ Phase 2 (2020) → to be replaced in Q2/2026

- ▶ 1022 standard nodes (2 Xeon Platinum 9242, 2x48 CPU cores, 384 GB memory)
- ▶ High Mem nodes: 16x 768 GB, 2x 1,5 TB
- ▶ 3 GPU nodes<sup>a</sup> (2 Xeon Gold 6148, 2x20 CPU cores, 768 GB memory, 4x NVIDIA V100)

## ■ Since 2021 operated for the NHR alliance



<sup>a</sup>The GPU nodes have been moved to Grete

# CPU System “Emmy”

Add-on procurement 2022: Phase 3

- NHR NEC CPU cluster (replacement of Emmy P1)
  - ▶ 2x Intel Sapphire Rapids 8468 (48 cores) per node
  - ▶ 164x 256 GB, 164x 512 GB, 12x 1 TB, 2x 2 TB nodes
- NHR+**SCC** NEC CPU add-on 2023
  - ▶ 20+44x 512 GB, 16+4x 1 TB, 1x2 TB
  - ▶ 4 nodes per 2U chassis, each with
    - 2x Intel Sapphire Rapids 8468 (48 cores) CPU
    - 1x Cornelis Omni-Path (100 Gbit/s) HCA



# GPU System “Grete”

- Technical specification: 103+3 nodes equipped with
  - ▶ 101 MEGWARE nodes (7 racks)
    - 2x AMD Epyc 7513 CPU (32 “Milan” cores, Zen 3 microarch.)
    - 512 GB memory (DDR4, 3200 MHz)
    - 2x 1 TB NVMe SSD
    - 4x NVIDIA A100 GPU (SXM4, 80/40 GB HBM2 memory)
    - 2x Mellanox InfiniBand HCA (HDR)
  - ▶ 2 nodes with twice the cores (Zen 2 microarch.), RAM, GPUs, and VRAM
  - ▶ 3 GPU nodes moved from Emmy Phase 2
- Installed at RZGö



# GPU System “Grete”

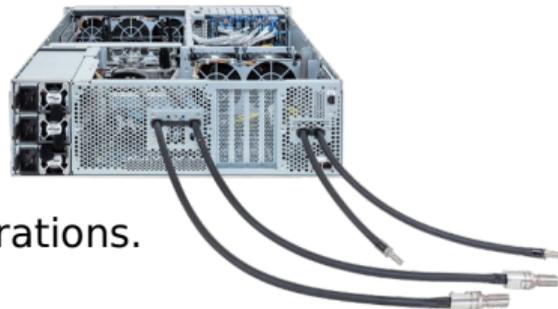
## ■ MEGWARE GPU cluster Grete

- ▶ Nodes from a range of projects (modular procurement):  
34x NHR, 22x REACT, 35x AI Service Center,  
9x Institutes/SFBs, 3x SCC

## ■ MEGWARE GPU Add-on 2024

- ▶ 25 nodes, each with
  - 2x Intel Sapphire Rapids 8468 (48 cores) CPU
  - 4x NVIDIA H100 SXM5 (94 GB) GPU
  - 2x InfiniBand HDR (200 Gbit/s) HCA
- ▶ 3U per node, DLC
- ▶ via RfP extension option AI Service Center

## ■ Joint operational concept since the start of operations.



# “Emmy” and “Grete” NG

Next generation 2026/2027: Phase 4

- New platform for the next CPU and GPU clusters at GWDG
- Modular procurement for NHR, SCC, HPC.NDS and KISSKI
- 350+ MEGWARE Eureka CPU nodes
- 2 installation phases
- AMD Turin and Venice CPUs (2x128 cores per node)
- 768-4096 GB memory
- 3,84 TB NVME SSDs
- Omni-Path CN5000 400G interconnect
- 15+ GPU nodes with 8xB200 Nvidia GPUs
- Reuse of CoolIT DLC infrastructure, racks and PDUs

# The “Future Technology Plattform” (FTP)

- Heterogeneous hardware for development tasks
- Direct access via shell and batch system
- Offered technologies:
  - ▶ Intel Habana Gaudi 2
  - ▶ NVIDIA Grace Hopper
  - ▶ GraphCore IPU
  - ▶ Esperanto.ai
  - ▶ SpiNNaker
  - ▶ Ampere Altra
  - ▶ NVIDIA Bluefield 2
  - ▶ AMD MI300A
- Full documentation and Quick Start Guide at [HPC docs](#)
- Courses on the heterogeneous hardware at the [GWDG Academy](#)

# Storage

- CephFS based HDD storage (23 PiB) successfully replaced old HDD Lustre as capacity storage
- New 1.6 PiB all-flash storage with MEGWARE supported Opensource Lustre as WORK storage at the MDC
- SSD based CephFS (0.75 PiB) well received für midrange performance applications
- VAST Data highly reliable enterprise storage with consistent low read latencies for AI
- DDN 2xES400NVX with Exa 6 and upgraded 15 TB SSD still sufficient fast main WORK storage at RZGÖ
- HSM storage succesfully migrated to new GWDG central HSM platform (currently still StorNext HSM, but probably Versity ScoutAM in the future)

## Selected Challenges Mentioned by Users

- Experience with Linux-Systems is limited
- Difficulty to debug distributed/parallel programs
- Getting access to the system
- Porting applications between centers
- Interactive access - replacing workstations
- Managing own projects, data and collaborators
- Ingress/Egress of data
- Need many concurrent users for training

# Means to Lower the Bar - Usability

- Users want to do science and don't want to become computer scientists
- Goal: Empower users to do science via
  - ▶ Tools
  - ▶ Prepared environments
  - ▶ Support
  - ▶ Consulting
  - ▶ Co-development
  - ▶ Training
  - ▶ Community
- Data management, data exploitation is part of science, too!

# JupyterHPC

- Quick and easy way to access HPC resources via
  - ▶ No ssh setup necessary
  - ▶ Frictionless and comfortable even for beginners
  - ▶ Very useful for courses, trainings and fallback for
- CPU and GPU nodes available
- Classic Jupyter notebooks with various kernels
  - ▶ Python, Julia, R, Bash, SageMath
- All container recipes open source
  - ▶ Easy to build and run your own containers

### Server Options

HPC Project (Username)  
Example project (u12345) Standard Profiles

HPC Type

Jupyter  Desktop

HPC Application

JupyterLab  RStudio  Code IDE

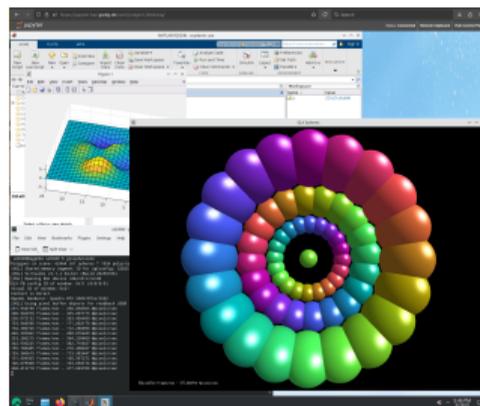
HPC Device

CPU  GPU

▶ Advanced

# HPC Desktops

- Goal: own HPC-powered workstation
- Integrated with Jupyter HPC
- Fully featured Linux desktops via webVNC
  - ▶ KDE, Gnome, XFCE available
- Fully hardware accelerated rendering on GPU nodes
  - ▶ Performant software rendering via Ilvmpipe on CPU nodes
- Complete HPC environment accessible
  - ▶ Regular module system, same as in standard ssh sessions
  - ▶ Documentation applies unchanged
  - ▶ Possible to start Slurm jobs from desktop sessions



## Data protection and security

- Secure HPC environment for batch processing of medical data
- Workflow adapted for other types of critical data
- Further improvements to simplify deployment and scale out in development
- Certification of processing data of highest non-classified data protection levels in preparation
- All OS image and configuration changes only with signed Git commits
- Critical backend system administration only by dedicated privileged admin workstations (PAW)
- Central logfile protection against manipulation
- Next step: Securing HPC deployment with Secure Boot (mostly done)

# Secure Boot

- Opensource deployment tool Warewulf
- Communication with Warewulf SSL secured
- Signed secure boot enabled iPXE bootloader
- Signed custom kernel modules with own custom key
- Deployed custom key to BIOS of all user accessible nodes
- Only one hardware platform did not work (soon to be phased out)

# What are privileged admin workstation

- Terminology from the Microsoft world
- Centrally managed devices for administration of critical infrastructure
- Maximize privilege separation (admin account of PAW has no admin privileges on the critical infrastructure and vice versa)
- Dedicated network segment (VPN profile)
- Restricted network access (no internet access)
- Restricted software availability (only tools necessary for the task)
- Logging of configuration changes and software installations

## Why and for what use PAW?

- In MS world: All tasks requiring domain admin privileges (root on all nodes of the domain)
- Jumphost for external admin access
- Central logserver
- Key management systems
- Build infrastructure for OS images
- Network management

# What is required and desired for setting up a PAW?

## ■ Required

- ▶ Automated deployment and configuration of PAW
- ▶ Monitoring of configuration and software installations on the PAW
- ▶ Central logging and analysis infrastructure for monitoring data
- ▶ Network or VPN profiles for accessing only the critical infrastructure via PAW
- ▶ Automated user setup

## ■ Desired

- ▶ No complicated device management infrastructure
- ▶ Personalized work environment on PAW should be possible
- ▶ Software installation from white listed pool

# GWDG HPC PAW solution

## ■ Management:

- ▶ Mainly based on Fedora and Kickstart
- ▶ Installation ISO and repo config from trusted mirror
- ▶ All configurations in internal Git repository
- ▶ Git repository requires signed commits
- ▶ List of approved keys for signature also in Git

## ■ Configuration

- ▶ UEFI Secure boot
- ▶ Systemd-boot without kernel commandline editor
- ▶ root account locked, LUKS, preconfigured sudo, firewalld, etc.
- ▶ Removal of remote management tools (SSH server, cockpit, etc.)
- ▶ pam-u2f for 2FA
- ▶ User setup includes eduroam and eduVPN setup
- ▶ Shell script for software installation kiosk
- ▶ Shell script with systemd timer for monitoring

# PAW logging infrastructure

- Graylog based
- Daily Systemd timer executes monitoring script and creates JSON
- Logging of Laptop type, Serial number, BIOS information
- rkhunter results
- All software packages with version (package manager and flatpaks)
- All executeables with set SUID, SGID or capabilities

## Why WORM logs?

Processing risk class D information (e.g. medical data) requires auditable access logs. This means:

- Central log aggregation
- Verifiable integrity
- Redundant long term storage
- Prevent changes to log files
- No deletion from a single person

## How to solve the requirements?

**Central:** Redirect Journald to Rsyslog, send Rsyslog to central server, individual files for each server

**Integrity:** Daily logfiles, create checksums

**Redundant storage:** Offsite backups (incl. checksums) on tape

**Deletion prevention** : root is god → difficult (can even change SELinux settings)

# File modification deletion prevention

Modification deletion prevention from normal users easy:

- Extended attributes append-only and immutable
- Creating new log files with append-only prevents deletion of older log entries
- Adding immutable attribute at the end of the day (before checksum creation) to prevent deletion
- Integrity check of scripts handling these operations with checksums

Superuser root still has permission to remove these extended attributes  
→ can modify or delete files and hide it by creating new checksums.

## Further integrity enhancing measurements

### Securing remote root logins:

- Changing the append-only and immutable attributes require `CAP_LINUX_IMMUTABLE`
- Removing this capability from `sshd` via `(immutable)` service override  
→ root via SSH has lost the permission to change append-only and immutable attribute
- Remote administration via SSH still possible
- Prevent usage of IPMI remote and serial console, e.g. no network connection
- Prevent local IPMI usage/configuration change by disabling OS access to BMC (BIOS setting)

## Further integrity enhancing measurements

### Securing local access:

- Server location in access restricted area of data center
- Persons with access to data center area must not know password for root or sudo enabled user accounts
- Persons with local login permissions must not have permission to access the restricted data center area alone
- Additionally require 2FA for local login (e.g. FIDO Key via pam-u2f)

# Trusted Research Environments (TREs)

- TREs provide **highly secure** environments for analyzing sensitive data
- All computations and analysis are conducted entirely within the TRE
  - ▶ Results are manually reviewed pre-publication
- Based on the "5 SAFE" framework from the UK Office of National Statistics
- Also known as:
  - ▶ **Secure Processing Environments (SPEs)** (Common in DE/EU)
  - ▶ Secure Data Environment (SDE) (Common in England; Especially NHS)
  - ▶ Data Safe Havens (Common in UK; Especially Scotland)

## Other Reasons for a TRE

### ■ Reproducibility

- ▶ Controlled Environments
- ▶ Container-Images are publishable

### ■ HPC-Resources (incl. GPU-Access)

- ▶ Big Data, for example Genome Analysis
- ▶ Machine-Learning Models

### ■ Outsourcing and Centralizing Security/Administration Burden

### ■ Legislative Pressure

- ▶ Europe: European Health Data Space
- ▶ Germany: Gesundheitsdatennutzungsgesetz

# 5 Safes

## 1 Safe People

- ▶ Trustworthy institutions and “Know your Customer”
- ▶ Completed training, Signed user agreement

## 2 Safe Projects

- ▶ No access without data owner's approval

## 3 Safe Settings

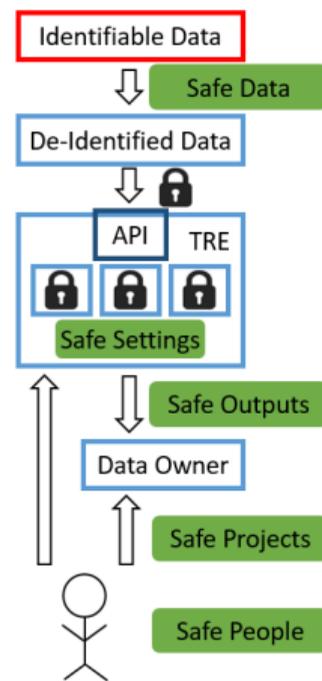
- ▶ Data is only processed within the TRE
- ▶ No internet access, no download possible

## 4 Safe Outputs

- ▶ Manual approval of exports before publication
- ▶ Focus: No re-identification; Only aggregated data

## 5 Safe Data

- ▶ Data is pseudonymized when ingested into the TRE



# On-Premise HPC-based TRE: GÖTRE

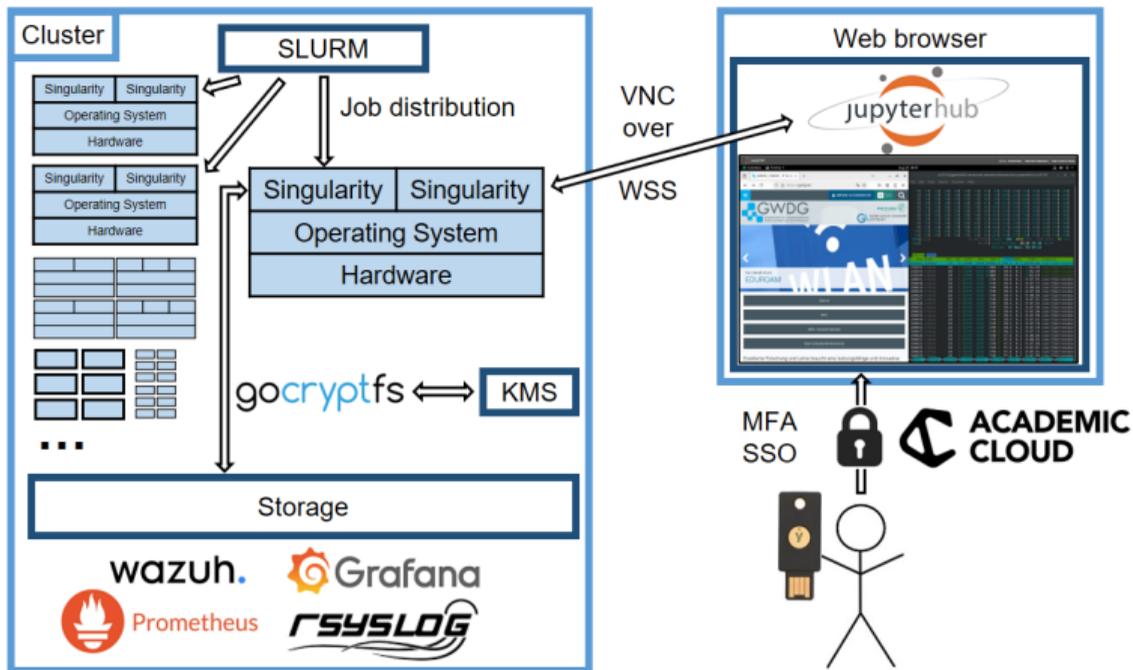
## Technical Features

- Zero-Trust Architecture
- End-to-End Encryption on file system level
  - ▶ Transparent to the user
- Multi-Factor Authentication
- Collaborative Work possible
  - ▶ Analogous to Windows Fileserver
- Complete Audit-Log
- One Key per File
  - ▶ Secure Sharing via Key Management System

## Organizational Features

- *High Usability despite Strong Security*
  - ▶ Access via web interface
  - ▶ No client software required
  - ▶ Graphical desktop environment
  - ▶ Supports non-web-based applications
- Embedded in HPC Infrastructure
  - ▶ High Performance
  - ▶ Machine-Learning support via GPUs

# GÖTRE: Architecture





# GöTRE: Interface

The screenshot displays a Jupyter Remote Desktop environment. The browser address bar shows `https://jupyter.hpc.gwdg.de/user/lars.quentin01/desktop/`. The Jupyter interface includes a top navigation bar with 'Status: Connected', 'Remote Clipboard', and 'Hub Control Panel'. The main workspace is split into two panes. The left pane shows a code editor with the following Python code:

```
[3]: from matplotlib import pyplot as plt
plt.subplot(1, 3, 1)
plt.title('First 9 images')
for i in range(9):
    plt.subplot(330 + 1 + i)
    plt.imshow(train_X[i], cmap=plt.get_cmap('gray'))
plt.show()
```

Below the code, a plot titled "First 9 images" displays a 3x3 grid of handwritten digits: 5, 0, 4, 1, 9, 2.

The right pane shows the output of the code, displaying training progress for epochs 4/15 through 12/15. Each line shows the training accuracy, validation accuracy, and validation loss. For example, at epoch 4/15, the training accuracy is 0.9893, validation accuracy is 0.9825, and validation loss is 0.0408.

# Outline

1 Site-update

2 Usability

3 Data protection and security

4 Conclusion

# Summary

- GWDG offers various HPC resources
- There are many ways to get access
- Besides performance, usability and availability can matter
- Providing highest data protection standards important esp. for medical research
- Combination of data protection and good usability possible